mathematics

MDPI

Article

Estimation-Robust Minimum-Variance Portfolio Allocation Under Factor-Loading Uncertainty

Leonardo P. Carvalho [1] **and Jonathan M. Palma** [2,*]

[1] Independent Researcher; eng.leonardocarvalho@gmail.com
[2] Faculty of Engineering, Universidad de Talca, Curicó 3340000, Chile
* Correspondence: jonathan.palma@utalca.cl

Abstract

We study sequential minimum-variance allocation when estimated market-factor loadings are uncertain. At each date, a horizon-one convex problem is solved using an ellipsoidal confidence region for the jointly estimated loading vector. The full loading covariance includes cross-asset residual dependence, and the uncertainty geometry is implemented either as a loading penalty or as an explicit worst-case exposure constraint. The principal walk-forward experiment uses the original nine U.S. sector exchange-traded funds from 2000 to 2024, yielding 6036 out-of-sample daily observations after a 252-day initial window. This fund-level opportunity set incorporates changes in underlying company constituents and therefore materially reduces the company-survivorship problem of a fixed stock panel, although the selected fund set remains subject to selection bias. An unrestricted past-only penalty selector over-regularizes: it chooses the largest normalized penalties on most dates and produces a negative long-short net return. We therefore report that frozen negative result and evaluate a separately frozen, economically guarded selector that caps the normalized penalty at 10. In long-short allocation, the guarded rule raises annualized net return from 1.13% under nominal beta targeting to 2.57%, lowers volatility from 10.72% to 8.39%, improves maximum drawdown from −38.67% to −24.42%, and reduces turnover and realized-beta dispersion. The paired Sharpe increase from 0.105 to 0.307 remains significant after false-discovery-rate correction. The result remains profitable across alternative estimation windows, transaction costs, and loading-covariance structures, and persists under a common monthly trading schedule. Its Sharpe is statistically indistinguishable from established uncertainty comparators, so universal dominance is not claimed. Long-only evidence favors zero or weak regularization. The results support economically bounded loading-uncertainty treatment in less-constrained portfolios, not unrestricted regularization or general robustness to distributional misspecification.

Keywords: estimation-error robust optimization; portfolio allocation; factor models; minimum-variance portfolios; beta targeting; transaction costs; ellipsoidal uncertainty; sequential optimization

MSC: 62P05

Check for updates

Academic Editors: Andrea Scozzari and Zengjing Chen

Received: 9 September 2026
Revised: 6 October 2026
Accepted: 12 October 2026
Published: 15 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Controlling exposure to systematic risk factors is a central problem in portfolio management [1–3]. The market beta plays a particularly important role: it directly determines the sensitivity of a portfolio to broad market movements. We emphasize at the outset that

the target beta $\beta^\star$ is treated here as an exogenous mandate parameter, a constraint the investor is required to satisfy, rather than a quantity chosen to maximize return. Two canonical mandates motivate this view. In a benchmark-relative or enhanced-index mandate, a long-only book must hold a prescribed market exposure (typically $\beta^\star = 1$) while improving on the benchmark along a secondary axis such as variance or turnover. In a market-neutral mandate, a long-short book must hold zero net market exposure ($\beta^\star = 0$) to target market neutrality at the estimated loadings; this does not guarantee zero realized market exposure or neutrality to other systematic factors. In neither case is the purpose of controlling beta to raise return, which it generally does not, but to deliver a contractually specified risk profile reliably. However, beta estimates are inherently noisy, being affected by sampling variability, structural breaks, and model misspecification [4,5]. Strategies that rely on point estimates of factor loadings are often unstable, leading to unintended exposures, excessive turnover, and degraded out-of-sample performance [6,7].

A second, well-documented difficulty in mean–variance optimization is that expected returns are substantially harder to estimate than covariances and estimation error can materially degrade out-of-sample allocations. This is not a claim that expected returns are universally "dominated by noise." Rather, it motivates minimum variance as a deliberately narrow objective and makes the naive $1/N$ rule an essential benchmark [8,9].

This paper studies sequential portfolio allocation under factor-loading estimation error. At each date, it solves a one-step problem using only trailing information and implements the resulting weights for the next return. We therefore use the precise term horizon-one receding-horizon optimization. The method does not optimize a multi-period control trajectory, and we do not claim a general model-predictive-control contribution.

Multi-period portfolio and MPC studies optimize sequences of controls or policies over a nontrivial horizon [10–12]. Our horizon-one implementation shares their repeated updating but not their multi-period trajectory optimization. Its closest methodological lineage is therefore sequential ellipsoidal robust portfolio optimization rather than a new MPC architecture.

The contribution has three separable parts: a portfolio design, an estimation-error treatment, and a sequential implementation. Specifically:

- The expected return term is removed from the objective, eliminating its associated estimation noise. In the long-short setting, where this removal would otherwise induce degeneracy, a single return-floor constraint preserves a non-trivial allocation.
- The target factor exposure is enforced as a hard equality constraint, avoiding a separate soft beta-targeting penalty and its tuning parameter.
- The factor-loading uncertainty set uses the joint OLS loading covariance, including cross-asset residual dependence, and its radius can be calibrated to a chi-square confidence level.
- Both the loading penalty and the explicit worst-case exposure constraint are evaluated; the former is a Lagrangian regularization, not the robust counterpart of the variance objective for an arbitrary fixed multiplier.
- The resulting optimization problem is a convex quadratic program with a single second-order cone term, and admits efficient solution via standard conic solvers.

The empirical contribution is a 2001–2024 walk-forward study on the original nine U.S. sector ETFs, supplemented by the submitted fixed-stock panel. It separates unrestricted from economically guarded loading regularization, reports the negative unrestricted calibration result rather than selecting it away, and evaluates alternative windows, costs, loading-covariance structures, trading frequencies, conventional benchmarks, and established uncertainty treatments.

Scope of the Contribution

The objective is reliable factor-exposure control under loading estimation error with limited trading. The revised evidence tests whether economically bounded regularization can also improve the realized risk–return trade-off relative to a matched nominal mandate; it does not assert universal dominance over established portfolio rules. The uncertainty set covers estimated loadings only and does not protect against misspecification of factor returns, residual distributions, volatility dynamics, tail dependence, or market regimes.

Terminology and Relation to Robustness Concepts

Four concepts that are sometimes grouped under "robust portfolio choice" must be separated. First, estimation-error robust optimization optimizes against parameter values in a finite-dimensional uncertainty set; the present loading ellipsoid and the framework of Goldfarb and Iyengar [13] belong here. Second, multiple-prior ambiguity aversion evaluates a decision under a set of priors; Garlappi, Uppal, and Wang [14] apply this idea to uncertain expected returns. Third, distributionally robust optimization optimizes over probability measures in a statistical-distance or moment set; Wasserstein methods are an example [15]. Fourth, Hansen–Sargent robust control allows an adverse distortion of a dynamic reference probability law, normally penalized by relative entropy and embedded in a Bellman problem [16].

Our procedure belongs to the first category. It contains no adversarial probability-law distortion, relative-entropy constraint, ambiguity-adjusted Bellman equation, or worst-case transition law. Accordingly, we remove robust-control claims and refer to the method as estimation-error robust minimum variance. The Garlappi–Uppal–Wang and Wasserstein implementations below are separate empirical comparators rather than alternative descriptions of our loading model.

The remainder of the paper is organized as follows. Section 2 introduces the factor model and nominal sequential problem. Section 3 derives the uncertainty set, penalty, and constrained formulation. Section 4 reports the walk-forward evidence and robustness checks. Section 5 concludes.

2. Model Formulation

2.1. Factor Model

We consider a universe of n risky assets whose returns follow a single-factor structure [7,17,18]. Let $R_t \in \mathbb{R}^n$ denote the vector of asset returns at time t, and let $f_t \in \mathbb{R}$ denote the factor return. The model is:

$$R_t = \alpha + \beta f_t + \epsilon_t, \tag{1}$$

where $\alpha \in \mathbb{R}^n$ is an intercept vector, $\beta \in \mathbb{R}^n$ is the vector of factor loadings, and ϵ_t is zero-mean noise with covariance Σ_ϵ.

Under this model, the asset-return covariance admits the decomposition:

$$\Sigma = \Sigma_\epsilon + \sigma_f^2 \beta\beta^\top, \tag{2}$$

where $\sigma_f^2 = \mathrm{Var}(f_t)$.

2.2. Portfolio Risk and Factor Exposure

For a portfolio with weights $w_t \in \mathbb{R}^n$, the portfolio variance and factor exposure are:

$$\mathrm{Var}(r_{p,t}) = w_t^\top \Sigma w_t, \qquad \beta_{p,t} = w_t^\top \beta. \tag{3}$$

We treat $\beta_{p,t}$ as a primary control variable: the investor specifies a target exposure $\beta^\star$, and the optimization is required to match it.

2.3. Nominal Minimum-Variance Objective

At each time step t, portfolio weights are obtained by minimizing a one-step objective that combines portfolio variance and a transaction cost term:

$$\min_{w_t \in \mathcal{W}} \quad w_t^\top \Sigma w_t \; + \; c \, \|w_t - w_{t-1}\|_1, \tag{4}$$

where $c > 0$ is the transaction cost coefficient and $\mathcal{W}$ is the feasible set. The ℓ_1 transaction-cost term models proportional execution costs and discourages large portfolio adjustments, promoting smoother allocation paths over time [19,20].

We do not include an expected-return term in the objective. The estimation of $\mathbb{E}[R]$ is known to be substantially more difficult than the estimation of Σ, and the resulting noise often dominates the benefit of return-seeking allocations in out-of-sample tests [8,9]. Removing this term eliminates a major source of instability while retaining a well-defined and computationally clean objective.

2.4. Beta Targeting as a Hard Constraint

In contrast to prior formulations that incorporate beta targeting via a soft quadratic penalty, we enforce the target exposure as a hard equality:

$$\hat{\beta}^\top w_t = \beta^\star, \tag{5}$$

where $\hat{\beta}$ is the OLS estimate of β obtained at time t. This change has two consequences. First, it removes the soft-penalty weight λ_β from the list of hyperparameters, simplifying calibration. Second, it ensures that the estimated portfolio beta matches the target whenever the constraint is feasible. Realized exposure can still deviate because of loading-estimation error and changes in the underlying loadings.

The constraint is feasible only when the asset universe and the remaining portfolio restrictions jointly span the target exposure. In the long-only simplex, a necessary and sufficient condition before adding further restrictions is $\beta^\star \in [\min_i \widehat{\beta}_i, \max_i \widehat{\beta}_i]$. We test this condition and the solver status at each rebalance rather than assuming feasibility.

It is worth making explicit the cost this constraint imposes in the minimum-variance setting, since it is easily mistaken for an inefficiency. Left unconstrained, a variance-minimizing long-only portfolio concentrates in low-volatility, defensive names and typically realizes a market beta well below one, which is a well-documented empirical property of minimum-variance portfolios.

Imposing $\hat{\beta}^\top w = \beta^\star = 1$ therefore moves the portfolio away from its unconstrained variance optimum and forgoes this low-beta tilt, raising realized variance relative to the unconstrained solution. The magnitude of this effect depends on the investable cross-section. It is muted on the large-cap universe studied here, where beta dispersion is limited (Section 4). But the qualitative point stands: this is not a flaw to be regularized away but the explicit price of the benchmark mandate. The constraint exists precisely because the investor is required to hold market exposure, and the framework's role is to meet that requirement at the least possible cost in variance and turnover.

In the market-neutral setting, the constraint plays the opposite, and arguably more essential, role. Here, $\beta^\star = 0$ is constitutive of the strategy rather than a tracking target, and removing it would change the very nature of the book. In the market-neutral setting, the constraint plays the opposite role, and arguably a more essential one. Here, $\beta^\star = 0$ is constitutive of the strategy rather than a tracking target. Removing it would change

the very nature of the book. Beta targeting is thus a risk-control and mandate-compliance device, not a return lever; the regimes differ only in whether the mandated exposure is the market or zero.

2.5. Feasible Sets and Portfolio Regimes

The framework accommodates different portfolio regimes through the choice of feasible set $\mathcal{W}$.

2.5.1. Long-Only Portfolios

The long-only setting is defined by:

$$\mathcal{W}_{LO} = \left\{ w \in \mathbb{R}^n \;\middle|\; \mathbf{1}^\top w = 1,\ w \geq 0 \right\}. \tag{6}$$

This corresponds to fully invested portfolios without short selling. The long-only constraint also acts as an implicit form of shrinkage in the spirit of Jagannathan and Ma [8].

2.5.2. Long-Short Portfolios

The market-neutral long-short setting is defined by:

$$\begin{aligned} \mathcal{W}_{LS} = \{ w \in \mathbb{R}^n \mid \mathbf{1}^\top w = 0,\ \|w\|_1 \leq L, \\ \hat{\mu}^\top w \geq \rho \}, \end{aligned} \tag{7}$$

where $L > 0$ is the leverage limit and $\rho > 0$ is a small return floor.

The return floor is a structural necessity in the long-short setting under a pure variance objective. Without it, the trivial solution $w = 0$ would minimize variance and satisfy all remaining constraints, producing a degenerate portfolio that never trades. The floor forces the optimizer to maintain a directional view sufficient to deliver at least ρ in expected return. We estimate $\hat{\mu}$ as the trailing window mean of asset returns, which is the simplest possible estimator. The loading-uncertainty treatment does not protect against estimation error in $\hat{\mu}$; the separate expected-return-uncertainty comparators are included precisely to keep that limitation visible.

2.6. Unified Optimization Problem

Combining the previous components, the nominal horizon-one problem at time t is:

$$\begin{aligned} \min_{w_t \in \mathcal{W}} \quad & w_t^\top \Sigma w_t \ + \ c \, \|w_t - w_{t-1}\|_1 \\ \text{s.t.} \quad & \hat{\beta}^\top w_t = \beta^\star, \end{aligned} \tag{8}$$

where $\mathcal{W} \in \{\mathcal{W}_{LO}, \mathcal{W}_{LS}\}$. This is a convex quadratic program in w_t and is solvable in milliseconds for moderate n using standard conic optimization methods.

3. Estimation-Robust Factor-Loading Treatment

3.1. Uncertainty in Factor Loadings

A key limitation of the nominal formulation (8) is the implicit assumption that $\hat{\beta}$ is known exactly. In practice, $\hat{\beta}$ is a noisy estimator, and the optimization treats it as if it were the truth. The hard beta-targeting constraint amplifies this issue: misestimation of $\hat{\beta}$ propagates one-for-one into errors in realized portfolio exposure.

Let the common regression design matrix be X and let $\widehat{\Sigma}_\epsilon$ denote the cross-sectional covariance matrix of the residual vector. All assets use the same aligned dates and the same intercept–market-factor design matrix. Consequently, $(X^\top X)^{-1}_{\hat{\beta},\hat{\beta}}$ below is the scalar

slope–slope element of the common 2×2 inverse design cross-product, not an asset-specific matrix. For the stacked OLS slope vector,

$$\widehat{\Sigma}_\beta = (X^\top X)^{-1}_{\beta,\beta} \widehat{\Sigma}_\epsilon. \tag{9}$$

Thus, $\mathrm{Cov}(\widehat{\beta}_i, \widehat{\beta}_j) = (X^\top X)^{-1}_{\beta,\beta} \widehat{\sigma}_{\epsilon,ij}$. The diagonal formula used in the submitted version is valid only after imposing zero cross-asset residual covariance. Our main specification uses the full matrix; diagonal and shrunk alternatives are reported as sensitivity checks.

For confidence level $1 - \alpha$, define $r_\alpha^2 = \chi^2_{n,1-\alpha}$ and

$$\mathcal{U}_\alpha = \left\{ \tilde{\beta} \in \mathbb{R}^n \;\middle|\; (\tilde{\beta} - \hat{\beta})^\top \widehat{\Sigma}_\beta^{-1} (\tilde{\beta} - \hat{\beta}) \le r_\alpha^2 \right\}. \tag{10}$$

This supplies an interpretable statistical radius rather than treating radius one as a 97-dimensional confidence set.

3.2. Robust Reformulation

We consider the worst-case factor exposure of a portfolio w over $\mathcal{U}$:

$$\sup_{\tilde{\beta} \in \mathcal{U}_\alpha} w^\top \tilde{\beta} = w^\top \hat{\beta} + r_\alpha \|Aw\|_2, \qquad A^\top A = \widehat{\Sigma}_\beta. \tag{11}$$

The closed-form follows from a standard duality argument for ellipsoidal robust optimization [21,22]. The second term represents the worst-case deviation in portfolio beta induced by estimation uncertainty.

The support term yields two distinct implementations. The first is a penalty motivated by worst-case exposure. The second directly constrains the absolute worst-case deviation from the mandate.

3.3. Penalty and Explicit Robust Constraint

The penalized problem is

$$\begin{aligned} \min_{w_t \in \mathcal{W}} \quad & w_t^\top \Sigma w_t + c\|w_t - w_{t-1}\|_1 + \gamma \|Aw_t\|_2 \\ \text{s.t.} \quad & \hat{\beta}^\top w_t = \beta^\star, \end{aligned} \tag{12}$$

where $A^\top A = \widehat{\Sigma}_\beta$. For an allowed worst-case exposure deviation δ, the explicit formulation instead imposes

$$|\widehat{\beta}^\top w_t - \beta^\star| + r_\alpha \|Aw_t\|_2 \le \delta. \tag{13}$$

Both formulations are convex second-order-cone programs. Under convexity, Slater regularity, and attainment, the constrained problem has a Lagrangian representation whose optimal multiplier supplies a corresponding penalty weight. This does not mean that an arbitrary fixed γ is equivalent to an arbitrary fixed δ, nor that every requested δ is feasible. We therefore report the penalty and constraint separately and record infeasible calibrated constraints explicitly.

3.4. Interpretation

The robust term admits several complementary interpretations.

Robust Optimization Perspective

$r_\alpha \|Aw\|_2$ is the support function governing the worst-case loading-induced beta perturbation. Penalization and constraint are related through Lagrangian duality only under the conditions stated above.

Regularization Perspective

From an optimization viewpoint, the robust term acts as a norm-based regularizer that penalizes exposure to directions in which factor uncertainty is high. This induces shrinkage toward more stable portfolios, conceptually similar to ridge regularization but tailored to the geometry of the estimator covariance.

Financial Interpretation

From a financial perspective, the robust term introduces a second layer of diversification. While Σ captures uncertainty in returns, Σ_β captures uncertainty in factor exposures. Penalizing $\|\Sigma_\beta^{1/2} w\|_2$ therefore discourages concentration in assets whose systematic risk is poorly estimated, even when their return statistics appear favorable.

3.5. Discussion

The extension preserves convexity and tractability. The multiplier γ is a tuning parameter; the confidence level $1 - \alpha$ controls the statistically calibrated ellipsoid; and δ is an exposure-tolerance parameter. Keeping these roles distinct is essential to interpretation.

Economically guarded calibration

A normalized penalty can still dominate the variance-and-cost objective if its admissible range is too broad. We therefore distinguish an unrestricted diagnostic grid from an economically guarded grid. The unrestricted grid is retained to reveal over-regularization. The guarded rule restricts $\widetilde{\gamma} \leq 10$, so the uncertainty term remains a stabilizer rather than becoming the de facto objective. This guard is evaluated as a second-stage exploratory redesign and is never described as part of the original prespecified result.

4. Empirical Study

4.1. Primary Sector-Fund Panel and Sample Limitations

The primary opportunity set comprises the original nine U.S. Select Sector SPDR funds: XLB, XLE, XLF, XLI, XLK, XLP, XLU, XLV, and XLY. Adjusted daily returns were downloaded using the Python `yfinance` library (version 1.2.0; https://pypi.org/project/yfinance/, accessed on 29 August 2026) and cover 3 January 2000–31 December 2024. After the initial 252-trading-day estimation window, the common out-of-sample period contains 6036 observations from 3 January 2001 through 31 December 2024. The market-regression factor is the S&P 500 price return. The long-only benchmarks are the S&P 500 Total Return index and adjusted SPY; the T-bill proxy is derived from `^IRX`.

The fund-level design materially reduces the company-survivorship problem raised by the fixed 97-stock panel: entry, exit, merger, and failure among the underlying companies are incorporated through the funds' realized returns rather than reconstructed from today's surviving firms. It does not eliminate selection bias. The nine funds are a fixed surviving set selected ex post, and the experiment is an asset-allocation study across sector funds rather than a point-in-time stock-index reconstruction. The submitted 97-stock experiment is retained only as a secondary diagnostic and is not the basis of the revised headline result.

4.2. Walk-Forward Design, Feasibility, and Inference

At date t, all parameters use observations through $t - 1$, the horizon-one problem is solved, and the weights are applied to return t. The central estimation window is 252 days, the transaction-cost coefficient is 0.002, and weights are updated daily. Long-only portfolios satisfy $\mathbf{1}^\top w = 1$, $w \geq 0$, and target beta one. Long-short portfolios satisfy $\mathbf{1}^\top w = 0$, target beta zero, gross leverage at most two, and the return-floor constraint stated in Section 2.

For long-only nominal targeting, feasibility of $\widehat{\beta}^{\top} w = 1$ on the simplex requires $1 \in [\min_i \widehat{\beta}_i, \max_i \widehat{\beta}_i]$ before additional constraints are considered. This convex-hull condition and solver status are checked at every rebalance. For the calibrated robust constraint, feasibility is determined jointly by the beta range, chi-square radius, uncertainty norm, and tolerance. Failures are recorded with their first date rather than dropped silently.

The conventional weight-based benchmarks—equal weight, inverse volatility, ERC, and the projected market-neutral ERC overlay—now receive the same rolling-beta and loading-uncertainty diagnostics as the optimized portfolios. For the S&P 500 Total Return and SPY return series, rolling realized beta is observable, but estimator-based worst-case portfolio exposure is not reported because constituent weights are unavailable.

Paired stationary-bootstrap inference uses 10,000 resamples and an expected block length of 21. Tests cover Sharpe ratio, volatility, maximum drawdown, turnover, rolling-beta dispersion, and worst-case exposure. Two-sided centered-bootstrap p-values are adjusted by the Benjamini–Hochberg procedure separately within each metric.

4.3. Penalty Calibration and Audit Sequence

The normalized coefficient entering the objective is

$$\gamma_t^{\text{eff}} = \widetilde{\gamma} \frac{s_{\Sigma,t}^2}{s_{\beta,t}}, \qquad s_{\Sigma,t}^2 = \frac{\operatorname{tr}(\Sigma_t)}{N^2}, \quad s_{\beta,t} = \frac{\sqrt{\operatorname{tr}(\Sigma_{\beta,t})}}{N}.$$

Every 63 trading days, a trailing 63-day inner walk-forward segment is formed inside the 252-day estimation window; candidate portfolios are re-estimated every five validation days. The original frozen grid was $\{0, 0.01, 0.03, 0.1, 0.3, 1, 3, 10, 30, 100\}$ and minimized mean squared net return. That loss emphasizes risk so strongly that the selector chose 100 on 5784 of 6036 long-short dates and generated a negative net return. We retain this result because it reveals a genuine over-regularization failure.

After diagnosing that failure, a separate exploratory manifest froze an economically guarded grid $\{0, 0.01, 0.03, 0.1, 0.3, 1, 3, 10\}$ before rerunning the strategy. The validation loss, timing, information set, and all other portfolio constraints were unchanged. Thus, the cap tests the economic proposition that loading uncertainty should stabilize rather than dominate the portfolio objective. It is reported transparently as a second-stage redesign, not as the original prespecified calibration.

4.4. Central Long-Only and Long-Short Results

Table 1 confirms that long-only is already strongly regularized by positivity, full investment, and the beta mandate. The least-intrusive selector chooses $\widetilde{\gamma} = 0$ throughout. The capped penalty marginally lowers volatility and reduces turnover, beta dispersion, and worst-case exposure, but its net Sharpe is 0.448 versus 0.456 for nominal beta targeting and 0.488 for ERC. We therefore do not recommend positive loading regularization as a return-enhancing long-only rule.

Beta dispersion is now reported for equal weight, inverse volatility, and equal-risk contribution using the same rolling factor regression applied to the optimized strategies. The external S&P 500 total-return and SPY series also receive realized-beta diagnostics. Their worst-case-beta entries remain not applicable because the ellipsoidal quantity requires a contemporaneous weight vector in the nine-fund opportunity set, which an external index return series does not provide.

The contrast in Table 2 is sharp. Nominal beta targeting earns 1.13% net with 10.72% volatility, a 0.105 Sharpe ratio, and a -38.67% maximum drawdown. The original unrestricted adaptive penalty suppresses volatility to 3.77% but also suppresses return below zero and raises turnover. The guarded selector instead earns 2.57% net, lowers volatility

to 8.39%, raises Sharpe to 0.307, improves drawdown to −24.42%, lowers turnover by 14%, and reduces rolling-beta dispersion by 23% relative to nominal. The paired Sharpe improvement remains significant after multiplicity correction ($p_{\mathrm{FDR}} = 0.0288$); the volatility, turnover, beta-dispersion, and worst-case-exposure improvements also survive correction.

Table 1. Central long-only sector-fund results. Return, volatility, and drawdown are percentages. Turnover is mean daily one-way absolute weight change. Worst beta is the 95th percentile of estimated worst-case deviation from the mandate.

Strategy	Net Return	Volatility	Sharpe	Max. DD	Turnover	Beta SD	Worst Beta
Nominal + beta	10.59	19.73	0.456	−58.05	0.00198	0.043	0.071
Capped adaptive penalty	10.37	19.60	0.448	−57.32	0.00149	0.035	0.055
Robust beta constraint	10.57	18.45	0.487	−53.01	0.00004	0.045	0.169
Equal weight	10.36	18.45	0.475	−52.53	0.00000	0.077	0.329
Inverse volatility	9.98	17.36	0.484	−49.18	0.00081	0.078	0.393
Equal risk contribution	9.98	17.20	0.488	−49.12	0.00106	0.085	0.414
S&P 500 total return	10.12	19.26	0.443	−55.25	0.00000	0.002	–
SPY adjusted	10.00	19.18	0.439	−55.19	0.00000	0.026	–

Table 2. Central long-short sector-fund results at native allocations. Return, volatility, and drawdown are percentages. All optimized portfolios use the same dollar-neutral mandate, leverage cap, return floor, and central transaction-cost coefficient.

Strategy	Net Return	Volatility	Sharpe	Max. DD	Turnover	Beta SD	Worst Beta
Nominal + beta	1.13	10.72	0.105	−38.67	0.01402	0.184	0.320
Original unrestricted adaptive	−0.40	3.77	−0.107	−17.84	0.01921	0.067	0.141
Capped adaptive penalty	2.57	8.39	0.307	−24.42	0.01200	0.142	0.277
Robust beta constraint	1.87	7.78	0.240	−29.47	0.00614	0.151	0.200
Goldfarb–Iyengar	1.42	6.55	0.217	−15.95	0.01035	0.106	0.190
Garlappi–Uppal–Wang	1.57	5.10	0.307	−12.39	0.01315	0.094	0.170
Market-neutral ERC	−4.01	11.97	−0.335	−74.73	0.01489	0.188	0.344

The central long-short risk–return comparison is summarized in Figure 1.

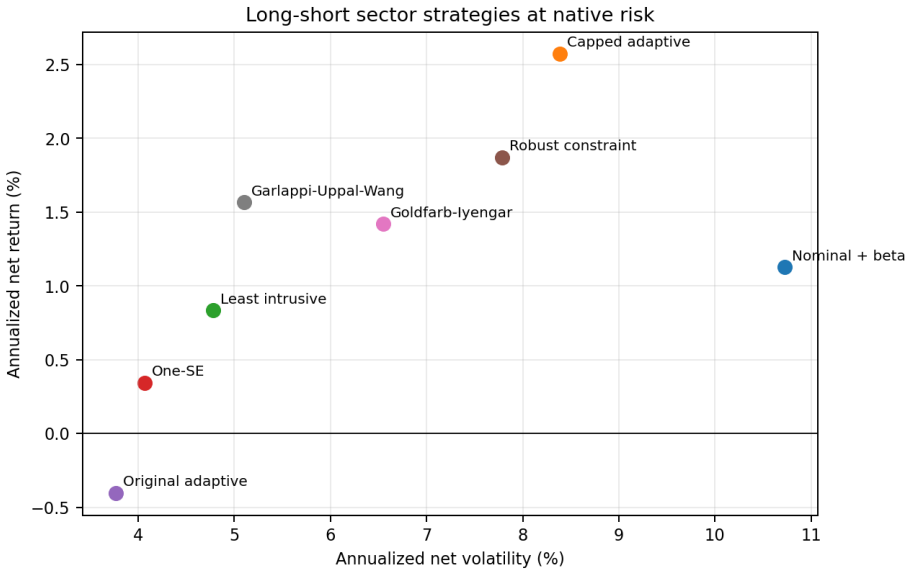


Figure 1. Annualized net return and volatility for the long-short sector strategies at native allocations. The original unrestricted adaptive selector is retained to show the consequence of over-regularization.

The guarded selector chooses $\widetilde{\gamma} = 10$ on 5217 of 6036 dates. Fixed-coefficient sensitivity explains the cap: values 10 and 30 retain positive return with Sharpe ratios of 0.294 and 0.273, respectively, whereas 100 pushes net return below zero. Moderate loading regularization is useful; unrestricted regularization is not.

4.5. Established Comparators and Feasibility

The Goldfarb–Iyengar-type comparator uses a jointly calibrated ellipsoid for expected-mean uncertainty, while the Garlappi–Uppal–Wang-type rule uses asset-specific multiple-prior confidence intervals. These are adaptations to the common rolling mandate rather than exact replications. The guarded selector has the highest point net return among the central robust formulations (2.57%), but Garlappi–Uppal–Wang has lower volatility (5.10%) and essentially the same point Sharpe (0.307). Pairwise Sharpe differences between the guarded rule and the three established uncertainty formulations are not statistically significant. The correct interpretation is competitiveness, not dominance.

The hard robust constraint also gives a balanced result: 1.87% net return, 7.78% volatility, 0.240 Sharpe, and lower turnover than the penalty. Across the complete sector calibration grid, 10 of 18 hard-constraint specifications solve, and 8 are infeasible. Long-short feasibility is only 3 of 9, and all successful long-short cases use tolerance 0.20. Infeasible specifications are excluded from conditional performance calculation by necessity, but not from the feasibility rate or narrative. The headline guarded-penalty comparison is unaffected because the penalized formulation remains feasible in all declared runs.

4.6. Targeted Robustness Wave

The guarded strategy remains profitable under every specification in Table 3. Results are stable across full, diagonal, and shrunk loading covariance. A 504-day window improves point Sharpe to 0.400, whereas a 126-day window weakens it to 0.093. A deliberately severe 50-basis-point transaction-cost coefficient reduces Sharpe to 0.077 but does not make net return negative. These weaker cases define important operational boundaries.

Table 3. Targeted long-short robustness checks for the capped adaptive rule. The monthly rows compare nominal and capped strategies under the same 21-day trading schedule.

Specification	Net Return	Volatility	Sharpe	Max. DD	Turnover	Beta SD
Central capped	2.57	8.39	0.307	−24.42	0.01200	0.142
Window 126	0.66	7.05	0.093	−29.69	0.01870	0.099
Window 504	3.64	9.10	0.400	−26.53	0.00737	0.147
Cost 0.0005	1.46	5.25	0.279	−11.33	0.01343	0.095
Cost 0.0010	2.71	6.58	0.411	−11.91	0.01153	0.113
Cost 0.0050	0.77	10.02	0.077	−35.26	0.01331	0.172
Covariance diagonal	2.80	8.59	0.326	−24.04	0.01198	0.151
Covariance shrunk	2.68	8.43	0.318	−24.21	0.01196	0.143
Monthly nominal	1.80	9.34	0.193	−30.15	0.00727	0.160
Monthly capped	2.19	7.93	0.276	−20.81	0.00652	0.141

Figure 2 presents the corresponding sensitivity of net return and Sharpe ratio across the targeted robustness specifications.

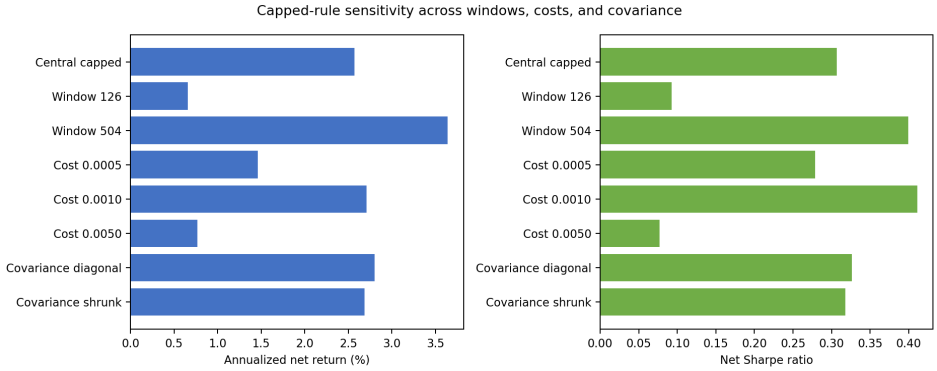


Figure 2. Net return and Sharpe sensitivity of the capped rule to estimation windows, transaction costs, and loading-covariance structures. The 252-day, full-covariance, 0.002-cost specification remains the declared central case.

Trading frequency does not explain the central result. With both strategies rebalanced every 21 days, the guarded rule earns 2.19% net versus 1.80% for nominal, lowers volatility from 9.34% to 7.93%, improves maximum drawdown from −30.15% to −20.81%, and lowers turnover. The volatility and turnover effects survive FDR correction; the monthly Sharpe increase from 0.193 to 0.276 is positive but not statistically significant.

4.7. External Validity and Remaining Limitations

The 24-year fund sample spans the technology bust, global financial crisis, pandemic sell-off, inflation shock, and multiple volatility regimes. Nevertheless, the nine funds are a selected surviving set, Yahoo Finance is not a licensed institutional database, and fund returns answer an asset-allocation question rather than a security-selection question. The sector design mitigates the specific company-survivorship problem but does not produce a fully selection-free experiment.

Market neutrality refers only to the estimated single market factor. Size, value, momentum, quality, and nonlinear exposures may remain. ETF-proxy factor regressions are available only over their common 2015–2024 period and are diagnostic rather than canonical Fama–French tests. Finally, the guarded cap was motivated by the frozen unrestricted failure and is therefore exploratory. Independent data are required for a clean confirmatory test of the guarded calibration.

5. Conclusions

This study characterizes the method narrowly as estimation-error robust minimum-variance allocation. It uses an ellipsoidal set for jointly estimated factor loadings and implements that geometry either as a loading penalty or as an explicit worst-case exposure constraint. It is not Hansen–Sargent robust control, does not optimize a multi-period trajectory, and does not claim protection against general distributional misspecification.

The primary evidence now comes from 6036 out-of-sample observations on the original nine U.S. sector funds. It reveals both a failure and a useful remedy. An unrestricted normalized-penalty grid over-regularizes the long-short allocation and produces a negative net return. A separately frozen guarded grid caps the normalized coefficient at 10. Relative to matched nominal beta targeting, that rule raises annualized net return from 1.13% to 2.57%, lowers volatility from 10.72% to 8.39%, improves maximum drawdown from −38.67% to −24.42%, reduces turnover and beta dispersion, and raises Sharpe from 0.105 to 0.307. The paired Sharpe improvement remains significant after FDR correction. The result remains profitable across alternative windows, transaction costs, covariance structures, and a common monthly trading schedule.

The guarded rule is competitive rather than universally superior. Its point Sharpe is essentially equal to the Garlappi–Uppal–Wang comparator, which achieves lower volatility, and Sharpe differences among the established uncertainty formulations are not significant. The explicit robust constraint offers lower turnover but is infeasible under tighter calibrations. In long-only allocation, the least-intrusive rule selects zero throughout, confirming that positivity, full investment, and the beta mandate already regularize the problem.

The fund panel incorporates changes among underlying companies and materially mitigates the fixed-stock survivorship objection, but the selected fund set remains ex post and surviving. The guarded cap was also motivated after observing the unrestricted failure. Accordingly, the contribution is evidence that economically bounded loading-uncertainty treatment can improve a matched long-short mandate over a long sector-fund sample, not a confirmatory claim of universal investment superiority.

Future work should test the guarded calibration on independent post-design data, use licensed point-in-time constituent histories where security selection is studied, and place genuine multifactor loading vectors and their joint uncertainty inside the optimizer. A multi-period extension would require optimizing a trajectory of future controls and states; a Hansen–Sargent extension would require an adversarial distortion of a dynamic probability law. Both remain distinct research programs.

Author Contributions: Conceptualization, L.P.C. and J.M.P.; methodology, L.P.C.; software, L.P.C.; validation, L.P.C. and J.M.P.; formal analysis, L.P.C.; investigation, L.P.C.; resources, J.M.P.; data curation, L.P.C.; writing—original draft preparation, L.P.C. and J.M.P.; writing—review and editing, L.P.C. and J.M.P.; visualization, L.P.C.; supervision, J.M.P.; project administration, J.M.P.; funding acquisition, J.M.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Agencia Nacional de Investigación y Desarrollo (ANID), Chile, through FONDECYT Regular Project N° 1241305.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All market series used in the study, including sector-fund returns, market benchmarks, and the T-bill proxy, were acquired from Yahoo Finance through the Python `yfinance` library (version 1.2.0; https://pypi.org/project/yfinance/, accessed on 29 August 2026). The repository records tickers, query dates, adjustment settings, raw downloads, derived returns, frozen experiment manifests with SHA-256 digests, successful and infeasible run registries, and scripts used to reproduce every reported table and figure. Yahoo Finance does not supply point-in-time S&P 100 membership; the fixed 97-stock panel is therefore secondary and explicitly labeled as survivorship-prone. Derived data and code are available from the corresponding author upon reasonable request.

Acknowledgments: The authors acknowledge support from FONDECYT Regular Project N° 1241305.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

1. Markowitz, H. The utility of wealth. *J. Political Econ.* **1952**, *60*, 151–158. [CrossRef]
2. Sharpe, W.F. Capital asset prices: A theory of market equilibrium under conditions of risk. *J. Finance* **1964**, *19*, 425–442. [CrossRef]
3. Rubinstein, M. Markowitz's "portfolio selection": A fifty-year retrospective. *J. Finance* **2002**, *57*, 1041–1045. [CrossRef]
4. Baesel, J.B. On the assessment of risk: Some further considerations. *J. Finance* **1974**, *29*, 1491–1494. [CrossRef]
5. Campa, J.M.; Goldberg, L.S. Employment versus wage adjustment and the US dollar. *Rev. Econ. Stat.* **2001**, *83*, 477–489. [CrossRef]
6. Michaud, R.O.; Ma, T. *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*; Oxford University Press: Oxford, UK, 2001.

7. Meucci, A. *Risk and Asset Allocation*; Springer: Berlin/Heidelberg, Germany, 2005.
8. Jagannathan, R.; Ma, T. Risk reduction in large portfolios: Why imposing the wrong constraints helps. *J. Finance* **2003**, *58*, 1651–1683. [CrossRef]
9. DeMiguel, V.; Garlappi, L.; Uppal, R. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Rev. Financ. Stud.* **2009**, *22*, 1915–1953. [CrossRef]
10. Li, D.; Ng, W.L. Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Math. Finance* **2000**, *10*, 387–406. [CrossRef]
11. Cadenillas, A.; Lakner, P.; Pinedo, M. Optimal control of a mean-reverting inventory. *Oper. Res.* **2010**, *58*, 1697–1710. [CrossRef]
12. Herzog, F.; Dondi, G.; Geering, H.P. Stochastic Model Predictive Control and Portfolio Optimization. *Int. J. Theor. Appl. Finance* **2007**, *10*, 203–233. [CrossRef]
13. Goldfarb, D.; Iyengar, G. Robust Portfolio Selection Problems. *Math. Oper. Res.* **2003**, *28*, 1–38. [CrossRef]
14. Garlappi, L.; Uppal, R.; Wang, T. Portfolio Selection with Parameter and Model Uncertainty: A Multi-Prior Approach. *Rev. Financ. Stud.* **2007**, *20*, 41–81. [CrossRef]
15. Esfahani, P.M.; Kuhn, D. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Math. Program.* **2018**, *171*, 115–166. [CrossRef]
16. Hansen, L.P.; Sargent, T.J. *Robustness*; Princeton University Press: Princeton, NJ, USA, 2008. [CrossRef]
17. Sharpe, W.F. A simplified model for portfolio analysis. *Manag. Sci.* **1963**, *9*, 277–293. [CrossRef]
18. Cochrane, J.H. *Asset Pricing*, revised ed.; Princeton University Press: Princeton, NJ, USA, 2005.
19. Gârleanu, N.; Pedersen, L.H. Dynamic trading with predictable returns and transaction costs. *J. Finance* **2013**, *68*, 2309–2340. [CrossRef]
20. Boyd, S.; Diamond, S.; Koh, K.; Nystrup, P.; Busseti, E.; Kahn, R.N.; Speth, J. Multi-period trading via convex optimization. *Found. Trends Optim.* **2017**, *3*, 1–76. [CrossRef]
21. Ben-Tal, A.; Nemirovski, A. *Robust Optimization–Methodology and Applications*; Springer: Berlin/Heidelberg, Germany, 2002; Volume 92, pp. 453–480.
22. Wang, H.; Zhao, Y.; Guo, Y.; Liu, C.; Zhang, X. Symmetric Adjustable Tail-Risk Measure for Distributionally Robust Optimization in Portfolio Allocation. *Symmetry* **2025**, *17*, 959. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.